



UNITED STATES PATENT AND TRADEMARK OFFICE

MN
UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/672,423	09/26/2003	John George Aschoff	IBMS.069PA(0510)	8426

62627 7590 04/12/2007
DAVID W. LYNCH
CHAMBLISS, BAHNER & STOPHEL
1000 TALLAN SQUARE-S
TWO UNION SQUARE
CHATTANOOGA, TN 37402

EXAMINER

TSAI, SHENG JEN

ART UNIT	PAPER NUMBER
----------	--------------

2186

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
2 MONTHS	04/12/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

**BEFORE THE BOARD OF PATENT APPEALS
AND INTERFERENCES**

Application Number: 10/672,423
Filing Date: September 26, 2003
Appellant(s): ASCHOFF ET AL.

MAILED

APR 12 2007

Technology Center 2100

David W. Lynch
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 01/05/2007 appealing from the Office action
mailed 06/26/2006.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal identifies the ground of rejections and the associated claims under rejection to be reviewed on appeal.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

6,820,035	Zahavi	11-16-2004
2003/0120864	Lee et al.	6-26-2003

Art Unit: 2186

(9) Grounds of Rejection

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1-9, 12-21, 24-33, and 36-40 are rejected under 35 U.S.C. 102(e) as being anticipated by Zahavi (US 6,820,035).

As to claim 1, Zahavi discloses **an administration device for providing automatic performance optimization of virtualized storage allocation within a network of storage elements** [System and Method for Determining Workload Characteristics for One or More Applications Operating in a Data Storage Environment (title); this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (abstract); figure 1], **comprising: memory for storing data thereon** [figure 1 shows a system memory (114) and a plurality of storage devices (115, 116) for storing data]; **and a processor** [the COMPUTER SYSTEM, figure 1, 113; In another embodiment, a program product includes a computer-readable medium having code included on the

Art Unit: 2186

medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, lines 66-67; column 3, lines 1-4)] **configured for receiving from a user a request for storage of data** [the user may request the type of disk drives (figure 6, step 184), the type of protection (figure 6, 186) for storage of data], **for determining workload requirements of the user making the request** [System and Method for Determining Workload Characteristics for One or More Applications Operating in a Data Storage Environment (title); this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (column 1, lines 23-28)], **for analyzing system parameters including performance characteristics** [ANALYZER API (figure 21, 600); ANALYZER ARCHIVES (figure 21, 602); adjustable performance comfort zone value (figure 6, 188); PERFORMANCE ZONE (figures 15-16); a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, 60-65); the analysis correlates all the selected devices according to system parameters (i.e., a matrix of coefficients) (figure 28, 838); The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in

a manner that meets a desired quality of service (column 7, lines 37-40); Summarized information may include data type, traffic requirements such as high as per second and performance characteristics such as MB per second for both front end (FE) and back end (BE), which are characteristics of the preferred data storage system and EMC Symmetrix (column 7, 49-53)] **of storage volumes within the network** [figure 1 shows the network (112) of a plurality of storage volumes (115, 116); figure 21, 614 shows the storage environment may be a SAN (Storage Area Network); figures 13-14, 21 and 23-27 illustrate analyzing system parameters and performance characteristics] **and for providing storage to meet the workload requirements of the user determined by the processor** [figure 10, step 224 shows the processor calculates workload requirements based on user provided request (steps 220 and 222); Determining the size and number of disk array or other data storage system needed by a customer requires information about both space, traffic and a desired quality of service. It is not sufficient to size a solution simply based on the perceived quantity of capacity desired, such as the number of terabytes believed to be adequate (column 1, lines 66-67; column 2, lines 1-4); The invention uses such information to advise a user on how to configure data storage systems having good capabilities to meet his needs, and while also considering traffic, other workload characteristics, and user defined Performance Zone Values. The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 33-40); In another embodiment, a program product includes a computer-readable

medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, lines 66-67; column 3, lines 1-4)] **and to meet competing workload requirements based on the analysis of the system parameters** [The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 37-40); a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, 60-65); ANALYZER API (figure 21, 600); ANALYZER ARCHIVES (figure 21, 602); the analysis correlates all the selected devices according to system parameters (i.e., a matrix of coefficients) (figure 28, 838); a "profile" of system parameters is generated by using a modeling tool (figure 29, 844); column 10, lines 32-48].

As to claim 2, Zahavi teaches that **the processor provides storage to meet the workload requirements based on the workload requirements of the user and storage requirements for the data** [The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 37-40); a program product includes a computer-readable

medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, 60-65); Determining the size and number of disk array or other data storage system needed by a customer requires information about both space, traffic and a desired quality of service. It is not sufficient to size a solution simply based on the perceived quantity of capacity desired, such as the number of terabytes believed to be adequate (column 1, lines 66-67; column 2, lines 1-4)].

As to claim 3, Zahavi teaches that **the processor provides storage to meet the workload requirements by selecting storage locations that meet performance and space requirements of the request** [The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 37-40); Determining the size and number of disk array or other data storage system needed by a customer requires information about both space, traffic and a desired quality of service. It is not sufficient to size a solution simply based on the perceived quantity of capacity desired, such as the number of terabytes believed to be adequate (column 1, lines 66-67; column 2, lines 1-4); these results take into consideration the performance constraints of the various components of the Symmetrix within each family, calculate the number of components requires and determine the architecture of each preferred Symmetrix model and build the required number of machines for the prescribed workload. In addition, the user is able to modify

the count of ports and directors in order to accommodate other needs such as redundancy, future growth or to account for uneven distribution of work, wherein the Logic compensates by calculating a totally balanced system. The number of back-end directors is calculated based on the number of required disks. There are physical limits to the number of disks that a back-end port can accommodate, depending on the preferred Symmetrix model. Sizing here is based on the maximum number of disks allowed per port (column 10, lines 32-48); storage space adjusted (figure 7, 198); figure 10].

As to claim 4, Zahavi teaches that **the processor selects storage locations that meet the performance and space requirements through analysis of the request for storage** [a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, 60-65); ANALYZER API (figure 21, 600); ANALYZER ARCHIVES (figure 21, 602); the analysis correlates all the selected devices according to system parameters (i.e., a matrix of coefficients) (figure 28, 838); a "profile" of system parameters is generated by using a modeling tool (figure 29, 844); column 10, lines 32-48; storage space adjusted (figure 7, 198); figure 10].

As to claim 5, Zahavi teaches that **the processor selects storage locations that meet the performance and space requirements through a storage policy mechanism** [Determining the size and number of disk array or other data storage

system needed by a customer requires information about both space, traffic and a desired quality of service. It is not sufficient to size a solution simply based on the perceived quantity of capacity desired, such as the number of terabytes believed to be adequate (column 1, lines 66-67; column 2, lines 1-4); The invention uses such information to advise a user on how to configure data storage systems having good capabilities to meet his needs, and while also considering traffic, other workload characteristics, and user defined Performance Zone Values. The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 33-40); a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, 60-65); ANALYZER API (figure 21, 600); ANALYZER ARCHIVES (figure 21, 602); the analysis correlates all the selected devices according to system parameters (i.e., a matrix of coefficients) (figure 28, 838); a "profile" of system parameters is generated by using a modeling tool (figure 29, 844); column 10, lines 32-48; storage policy mechanism includes "type of disk drives" (figure 6, 184), "type of protection (e.g., RAID-1, RAID-S, unprotected)" (figure 6, 186)].

As to claim 6, Zahavi teaches that **the processor determines workload attributes of the user** [this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for

determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (abstract); specifically, this invention is directed to a configuration method and system for storage capacity planning based on user or administrator defined workload requirements (abstract); figure 13, 524 shows an user interface that allows a user to define workload requirement; figure 14] **and desired levels of performance** [adjustable performance comfort zone value (figure 6, 188); PERFORMANCE ZONE (figures 15-16)] , **retains the latest information about the available capacity within the network of storage elements** [configure storage network, figure 12; configure other data storage systems , figure 11], **determines performance characteristics of individual storage devices at different locations within the network as a function of the workload requirements of the user** [figures 6-21], **and determines a presence and attributes of competing workloads sharing the storage devices over extended periods of time** [figures 22-27].

As to claim 7, Zahavi teaches that **the processor is configured for determining workload requirements of the user** [System and Method for Determining Workload Characteristics for One or More Applications Operating in a Data Storage Environment (title); this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (column 1, lines 23-28)] **by using canned workload descriptions that are based on**

characterizations of user environments across various industries and applications [The corresponding canned workload description in Zahavi's invention is the Library of Workload shown in figure 5, 172; figure 13, 524 shows an user interface that allows a user to define workload requirement; figure 14; "type of disk drives" (figure 6, 184), "type of protection (e.g., RAID-1, RAID-S, unprotected)" (figure 6, 186)].

As to claim 8, Zahavi teaches that **the processor is configured for determining workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user** [System and Method for Determining Workload Characteristics for One or More Applications Operating in a Data Storage Environment (title); this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (column 1, lines 23-28); The Workload Profiler (figure 21,610) is an automated tool that correlates workload data on a storage system, as collected and saved by analyzing tool such as the ECC Workload Analyzer from EMC Corporation of Hopkinton, Mass. The workload data may be IO data, or response time, on any other metric useful for measuring work in a storage environment. IO traffic is a basic indicator of work for a logical device, and response times may be used as a metric for identifying specific problem spots and correlate these problem spots with applications thus showing root-cause to application response time problems (column 12, lines 4-14)].

As to claim 9, Zahavi teaches that **the processor is configured for determining workload requirements of the user** [System and Method for Determining Workload Characteristics for One or More Applications Operating in a Data Storage Environment (title); this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (column 1, lines 23-28)] **by using intelligent software components that analyze workload descriptions for an application of the user** [a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, 60-65); ANALYZER API (figure 21, 600); ANALYZER ARCHIVES (figure 21, 602); the analysis correlates all the selected devices according to system parameters (i.e., a matrix of coefficients) (figure 28, 838)].

As to claim 12, Zahavi teaches that **the processor is configured for performing a calibration process to discover the performance capabilities of the underlying storage devices** [figure 3 shows the performance degradation factor; figures 5-9 show the process of adjusting storage allocation based on the type of disks and their characteristics].

As to claim 13, Zahavi discloses **a network storage system** [FIG. 1 is a block diagram of a data storage network for which Logic (FIG. 2) that is part of the computer system shown in FIG. 1 is particularly useful (column 13, lines 11-13)], **comprising:**

a plurality of storage devices [figure 1 shows a plurality of storage devices (115 and 116)];

a plurality of servers [figure 1 shows a service processor (123); in addition, remote systems (111), which have their own respective service processor, are to be configured into the system as well, resulting a plurality of service processors] **coupled to the plurality-of storage devices** [figure 1], **via network interconnect** [remote systems connected via network]; **and**

an administration device, coupled to at least the plurality of storage devices, for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, wherein the administration device further comprises: memory for storing data thereon [refer to "As to claim 1" through "As to claim 9"]; **and**

a processor [the COMPUTER SYSTEM, figure 1, 113; In another embodiment, a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, lines 66-67; column 3, lines 1-4)] **configured for receiving from a user a request for storage of data** [the user may request the type of disk drives (figure 6, step 184), the type of protection (figure 6, 186)]

Art Unit: 2186

for storage of data], **for determining workload requirements of the user making the request** [System and Method for Determining Workload Characteristics for One or More Applications Operating in a Data Storage Environment (title); this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (column 1, lines 23-28)], **for analyzing system parameters including performance characteristics** [System and Method for Determining Workload Characteristics for One or More Applications Operating in a Data Storage Environment (title); this invention relates generally to managing and analyzing data in a data storage environment, and more particularly to a system and method for determining workload characteristics including the profiles for such characteristics for one or more applications operating in a data storage environment (column 1, lines 23-28); The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 37-40); Summarized information may include data type, traffic requirements such as high as per second and performance characteristics such as MB per second for both front end (FE) and back end (BE), which are characteristics of the preferred data storage system and EMC Symmetrix (column 7, 49-53)] **of storage volumes within the network** [figure 1 shows the network (112) of a plurality of storage volumes (115, 116); figure 21, 614 shows the storage environment may be a SAN (Storage Area

Art Unit: 2186

Network); figures 13-14, 21 and 23-27 illustrate analyzing system parameters and performance characteristics] **and for providing storage to meet the workload requirements of the user determined by the processor** [figure 10, step 224 shows the processor calculates workload requirements based on user provided request (steps 220 and 222); Determining the size and number of disk array or other data storage system needed by a customer requires information about both space, traffic and a desired quality of service. It is not sufficient to size a solution simply based on the perceived quantity of capacity desired, such as the number of terabytes believed to be adequate (column 1, lines 66-67; column 2, lines 1-4); The invention uses such information to advise a user on how to configure data storage systems having good capabilities to meet his needs, and while also considering traffic, other workload characteristics, and user defined Performance Zone Values. The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 33-40); In another embodiment, a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, lines 66-67; column 3, lines 1-4)] **and to meet competing workload requirements based on the analysis of the system parameters** [a program product includes a computer-readable medium having code included on the medium configured to carry out computer-executed steps of analyzing

Art Unit: 2186

work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, 60-65); ANALYZER API (figure 21, 600); ANALYZER ARCHIVES (figure 21, 602); the analysis correlates all the selected devices according to system parameters (i.e., a matrix of coefficients) (figure 28, 838); a "profile" of system parameters is generated by using a modeling tool (figure 29, 844); column 10, lines 32-48].

As to claim 14, refer to "As to claim 2."

As to claim 15, refer to "As to claim 3."

As to claim 16, refer to "As to claim 4."

As to claim 17, refer to "As to claim 5."

As to claim 18, refer to "As to claim 6."

As to claim 19, refer to "As to claim 7."

As to claim 20, refer to "As to claim 8."

As to claim 21, refer to "As to claim 9."

As to claim 24, refer to "As to claim 12."

As to claim 25, refer to "As to claim 1."

As to claim 26, refer to "As to claim 2."

As to claim 27, refer to "As to claim 3."

As to claim 28, refer to "As to claim 4."

As to claim 29, refer to "As to claim 5."

As to claim 30, refer to "As to claim 6."

As to claim 31, refer to "As to claim 7."

Art Unit: 2186

As to claim 32, refer to "As to claim 8."

As to claim 33, refer to "As to claim 9."

As to claim 36, refer to "As to claim 12."

As to claim 37, refer to "As to claim 1."

As to claim 38, refer to "As to claim 6."

As to claim 39, refer to "As to claim 1."

As to claim 40, refer to "As to claim 1."

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 10-11, 22-23, and 34-35 are rejected under 35 U.S.C. 103(a) as being unpatentable over Zahavi (US 6,820,035), and in view of Lee et al. (US Patent Application Publication 2003/0120864).

As to claim 10, Zahavi does not explicitly mention that **the processor is configured for accessing a virtualization engine and volume managers to stripe data within a virtual disk across managed storage devices.**

However, the apparatus disclosed by Zahavi is a generic storage network system that is directly applicable to support storage virtualization.

Further, Lee et al. teach in their invention "High-Performance Log-Structured RAID" a method and apparatus of improving the performance of a RAID system under

particular workload condition [paragraph 0007] for supporting storage virtualization [abstract; paragraph 0002] in which stripes of data are distributed among a plurality of disks [figure 2; paragraph 0096].

Storage virtualization allows dynamic allocation of data based on the availability of each storage device in the system, which results in better utilization of the capacity of the storage devices. Stripe of data across storage devices provides redundancy and allows data to be recovered upon failures, hence improves the reliability of the system.

Therefore, it would have been obvious for one of ordinary skills in the art to recognize the benefits of employing storage virtualization and data striping, as demonstrated by Lee et al., and to incorporate it into the existing apparatus disclosed by Zahavi to fully utilize the capacity of the storage devices and to improve the reliability of the system.

As to claim 11, Lee et al. teach that **the processor is configured for determining how to relocate virtual disks to meet a desired level of performance** [stripes of data are distributed among a plurality of disks (figure 2; paragraph 0096) to improve reliability of the storage system].

As to claim 22, refer to "As to claim 10."

As to claim 23, refer to "As to claim 11."

As to claim 34, refer to "As to claim 10."

As to claim 35, refer to "As to claim 11."

(10) Response to Arguments

Applicants' arguments have been fully and carefully considered with Examiner's answers set forth below.

Response to Arguments on Independent Claims 1, 13, 39 and 40 (i)

Applicants contend that independent claims 1, 13, 39 and 40 are patentable over Zahavi (US 6,820,035) as, allegedly, Zahavi fails to teach, disclose or suggest a processor configured for determining workload requirements of a user making a request for storage of data, simply because Zahavi's invention allows a user provides certain information through the user interface. The Examiner disagrees with this argument for the following reasons:

First, Zahavi's invention is directed to an automatic tool which creates workload profiles based on collected data, and **the corresponding processor configured to determine workload profile is the Workload Profiler** [figure 21, 610; column 11, lines 40-44; column 12, lines 4-14]. Zahavi teaches the following regarding the Workload Profiler:

- figure 21, 610 shows an I/O Data Profile Creator that correlates, clusters (auto grouping), creates profiles and save profiles; figure 21, 612 shows "Built Profiles;"
- a Profile Creator or Profiler (figure 21,610) which receives analyzed IP data by communicating with analyzer application program interface (API) 600, which may receive data for analyzing from stored analyzer archives 602 (column 11, lines 40-44);

- the Workload Profiler (figure 21,610) is an automated tool that correlates workload data on a storage system, as collected and saved by analyzing tool such as the ECC Workload Analyzer from EMC Corporation of Hopkinton, Mass. The workload data may be IO data, or response time, on any other metric useful for measuring work in a storage environment. IO traffic is a basic indicator of work for a logical device, and response times may be used as a metric for identifying specific problem spots and correlate these problem spots with applications thus showing root-cause to application response time problems (column 12, lines 4-14);
- figure 29, step 844 shows the flow chart of the automatic process where a profile is generated for each of the affinity group.
- figures 13-16 show the automatic tool with a configuration tool (figure 13, 148) including a Workload Characterization unit (figure 13, 524 and figure 14);

Second, although Zahavi's invention allows a user to provide certain parameters to the Workload Profiler [column 7, lines 13-15], but the actual calculation and creation of workload profiles are performed by the automatic tool, i.e., the Workload Profiler [The logic of the invention calculates the amount of work done on the front-end and back-end of the data storage system, in this example, the preferred Symmetrix. Using these numbers together with configuration selections the number of Symmetrix is calculated and the number and type of front-end ports is presented (column 10, lines 23-28)].

Third, it is noted that having a user provide parameters to the system is consistent with Applicants' disclosure, because Paragraph 0045 of the Specification section of the Application [Aschoff et al., US Patent Application Publication 2006/0161753] states "A user interface, such as a graphic interface, cooperatively coupled to virtual disk allocator 272, allows a user to point to a grouping of volumes and a particular window of time, and then create a workload profiles based on the observed behavior of those volumes." Thus, Zahavi's scheme of how a workload profile is created is consistent with Applicants' disclosure, and teaches the limitation of the claims.

Fourth, the claim language recites "a processor configured for receiving from a user a request for storage of data, for determining workload requirements of the user making the request, ..." It is noted that user provided parameters associated with a request are considered as part of the request. Thus Zahavi's teaching is also consistent with the limitations recited in the claims.

Response to Arguments on Independent Claims 1, 13, 39 and 40 (ii)

Applicants contend that independent claims 1, 13, 39 and 40 are patentable over Zahavi (US 6,820,035) as, allegedly, Zahavi fails to teach, disclose or suggest a processor configured for providing storage to meet the workload requirements of the user determined by the processor and to meet competing workload requirements based on the analysis of the system parameters, simply because Zahavi's invention allows a user provides certain information through the user interface. The Examiner disagrees with this argument for the following reasons:

First, the allegation that the workload requirements are provided by the user instead of being created by the processor has been fully addressed in the section presented earlier in this Office Action [the corresponding processor configured to determine workload profile is the Workload Profiler (figure 21, 610; column 11, lines 40-44; column 12, lines 4-14), Refer to "*Response to Arguments on Independent Claims 1, 13, 39 and 40 (i)*" for details].

Second, as for the limitation that "a processor configured to providing storage to meet the workload requirements," Zahavi's teachings include:

- Determining the size and number of disk array or other data storage system needed by a customer requires information about both space, traffic and a desired quality of service. It is not sufficient to size a solution simply based on the perceived quantity of capacity desired, such as the number of terabytes believed to be adequate (column 1, lines 66-67; column 2, lines 1-4);
- The invention uses such information to advise a user on how to configure data storage systems having good capabilities to meet his needs, and while also considering traffic, other workload characteristics, and user defined Performance Zone Values. The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 33-40);
- In another embodiment, a program product includes a computer-readable medium having code included on the medium configured to carry out computer-

executed steps of analyzing work-related data for creating a correlation of logical devices and then using the correlation to perform a storage management function (column 2, lines 66-67; column 3, lines 1-4).

Therefore, the Examiner's position regarding the patentability of claims 1, 13, 39 and 40 remains the same as indicated in the previous Office Action.

Response to Arguments on Independent Claims 25 and 37

Applicants contend that independent claims 25 and 37 are patentable over Zahavi (US 6,820,035) as, allegedly, Zahavi fails to teach, disclose or suggest a processor configured for providing automatic performance optimization of virtualized storage allocation with a network of storage elements by determining workload requirements of a user making a request for storage of data, simply because Zahavi's invention allows a user provides certain information through the user interface. The Examiner disagrees with this argument for the following reasons:

First, the allegation that the workload requirements are provided by the user instead of being created by the processor has been fully addressed in the section presented earlier in this Office Action [**the corresponding processor configured to determine workload profile is the Workload Profiler** (figure 21, 610; column 11, lines 40-44; column 12, lines 4-14), Refer to "***Response to Arguments on Independent Claims 1, 13, 39 and 40 (i)***" for details].

Second, as for the limitation that "automatic performance optimization," Zahavi's teachings include:

- There is a long-felt need for a computer-based tool that would allow a straightforward non-complex way to allocate proper storage capacity while balancing cost, growth plans, workload, and performance requirements (column 2, lines 5-8);
- The invention allows for the integration of space and traffic needs of a business along with performance goals such that the resulting configuration can handle the workload in a manner that meets a desired quality of service (column 7, lines 37-40);
- Summarized information may include data type, traffic requirements such as high as per second and performance characteristics such as MB per second for both front end (FE) and back end (BE), which are characteristics of the preferred data storage system and EMC Symmetrix (column 7, 49-53);
- adjustable performance comfort zone value, figure 6, 188, column 6, lines 6-23;
- PERFORMANCE ZONE, figure 16.

Therefore, the Examiner's position regarding the patentability of claims 25 and 37 remains the same as indicated in the previous Office Action.

Response to Arguments on Dependent Claims 6 and 18

Applicants contend that dependent claims 6 and 18 are patentable over Zahavi (US 6,820,035) because they recite a processor that determines workload attributes of the user and desired levels of performance and, allegedly, Zahavi relies on user input for the workload attributes. The Examiner disagrees.

The allegation that the workload requirements are provided by the user instead of being created by the processor has been fully addressed in the section presented earlier in this Office Action [**the corresponding processor configured to determine workload profile is the Workload Profiler** (figure 21, 610; column 11, lines 40-44; column 12, lines 4-14), Refer to “***Response to Arguments on Independent Claims 1, 13, 39 and 40 (i)***” for details].

The issue of determining performance level also has been fully addressed in the section presented earlier in this Office Action. Refer to “***Response to Arguments on Independent Claims 25 and 37***” for details.

Therefore, the Examiner’s position regarding the patentability of claims 6 and 18 remains the same as indicated in the previous Office Action.

Response to Arguments on Dependent Claims 7 and 19

Applicants contend that dependent claims 7 and 19 are patentable over Zahavi (US 6,820,035) by alleging, again, that Zahavi relies on user input for the workload requirements. The Examiner disagrees for the following reasons:

First, the allegation that the workload requirements are provided by the user instead of being created by the processor has been fully addressed in the section presented earlier in this Office Action [**the corresponding processor configured to determine workload profile is the Workload Profiler** (figure 21, 610; column 11, lines 40-44; column 12, lines 4-14), Refer to “***Response to Arguments on Independent Claims 1, 13, 39 and 40 (i)***” for details].

Second, note that claims 7 and 19 recite “a processor that is configured for determining workload requirements of the user by using canned workload description that are based on characterization of user environments across various industries and applications.”

The corresponding canned workload description in Zahavi's invention is the Library of Workload shown in figure 5, 172.

Therefore, the Examiner's position regarding the patentability of claims 7 and 19 remains the same as indicated in the previous Office Action.

Response to Arguments on Dependent Claims 8 and 20

Applicants contend that dependent claims 8 and 20 are patentable over Zahavi (US 6,820,035) by alleging that Zahavi fails to mention the use of observations of storage access patterns of a user. The Examiner disagrees.

Zahavi's invention is directed to an automatic tool which creates workload profiles based on collected data, and **the corresponding processor configured to determine workload profile is the Workload Profiler** [figure 21, 610; column 11, lines 40-44; column 12, lines 4-14]. Particularly, Zahavi teaches

- A Profile Creator or Profiler (figure 21,610) which receives analyzed IP data by communicating with analyzer application program interface (API) 600, which may receive data for analyzing from stored analyzer archives 602 (column 11, lines 40-44);
- The Workload Profiler (figure 21,610) is an automated tool that correlates workload data on a storage system, as collected and saved by analyzing tool

such as the ECC Workload Analyzer from EMC Corporation of Hopkinton, Mass. The workload data may be IO data, or response time, on any other metric useful for measuring work in a storage environment. IO traffic is a basic indicator of work for a logical device, and response times may be used as a metric for identifying specific problem spots and correlate these problem spots with applications thus showing root-cause to application response time problems (column 12, lines 4-14).

Thus, storage access data is not only corrected but also analyzed in order to create a workload profile and achieve the desired performance level.

Therefore, the Examiner's position regarding the patentability of claims 8 and 20 remains the same as indicated in the previous Office Action.

Response to Arguments on Dependent Claims 9 and 21

Applicants contend that dependent claims 9 and 21 are patentable over Zahavi (US 6,820,035) by alleging that Zahavi fails to mention using intelligent software components that analyze workload description of an application. The Examiner disagrees.

Zahavi's invention is directed to an automatic tool which creates workload profiles based on collected data, and **the corresponding processor configured to determine workload profile is the Workload Profiler** [figure 21, 610; column 11, lines 40-44; column 12, lines 4-14]. Particularly, Zahavi teaches

- Analyzer API (figure 21, 600) and Analyzer Archives (figure 21, 602) that are used to create workload profile;

Art Unit: 2186

- A Profile Creator or Profiler (figure 21,610) which receives analyzed IP data by communicating with analyzer application program interface (API) 600, which may receive data for analyzing from stored analyzer archives 602 (column 11, lines 40-44);
- The Workload Profiler (figure 21,610) is an automated tool that correlates workload data on a storage system, as collected and saved by analyzing tool such as the ECC Workload Analyzer from EMC Corporation of Hopkinton, Mass. The workload data may be IO data, or response time, on any other metric useful for measuring work in a storage environment. IO traffic is a basic indicator of work for a logical device, and response times may be used as a metric for identifying specific problem spots and correlate these problem spots with applications thus showing root-cause to application response time problems (column 12, lines 4-14).

Thus, storage access data is not only corrected but also analyzed in order to create a workload profile and achieve the desired performance level.

Therefore, the Examiner's position regarding the patentability of claims 9 and 21 remains the same as indicated in the previous Office Action.

Response to Arguments on Dependent Claims 10-11, 22-23, 34-35 and 41-42

Applicants contend that dependent claims 10-11, 22-23, 34-35 and 41-42 are patentable over Zahavi (US 6,820,035) by virtue of inheriting the limitations of independent claims 1, 13 and 25. The Examiner disagrees.

As established in the section of "**Response to Arguments on Independent Claims 1, 13, 39 and 40 (i)**," the allegation that Zahavi fails to teach, disclose or suggest a processor configured for determining workload requirements of a user making a request for storage of data simply because Zahavi's invention allows a user provides certain information through the user interface is false **the corresponding processor configured to determine workload profile is the Workload Profiler** (figure 21, 610; column 11, lines 40-44; column 12, lines 4-14)]. Zahavi teaches the following regarding the Workload Profiler:, and claims 1, 13 and 25 remain rejected under 35 U.S.C. 102(e) as being anticipated by Zahavi (US 6,820,035).

Further, the reason why each of dependent claims 10-11, 22-23, 34-35 is rejected has been provided in section (9) of this Office Action (i.e., **(9) Grounds of Rejection**).

Therefore, the Examiner's position regarding the patentability of claims 10-11, 22-23, 34-35 remains the same as indicated in the previous Office Action.

(11) Related Proceedings Appendix

None.

Art Unit: 2186


Sheng-Jen Tsai
Examiner
Art Unit 2186

Sheng Jen Tsai



Matthew Kim
Supervisory Patent Examiner
Art Unit 2186

Lynne Browne
Supervisory Patent Examiner
Appeal Conference Specialist
Technology Center 2100


Lynne H. Browne
Appeal Specialist, TQAS
Technology Center 2100

April 10, 2007